

terms/phrases of the query for example, and the association is stored. (Block 520) Referring back to block 510, if an aggregation condition (e.g., the expiration of a time period, a command for aggregation, the acquisition of a certain amount of data, etc.) is met, previously stored document/domain information to query information associations are aggregated (Block 530) and the aggregated information is stored (Block 540). The method 500 may be left upon the occurrence of an exit command or condition. (Node 550)

[0060] FIGS. 6A and 6B illustrate exemplary data structures that may be used to store aggregated selected document/domain information to query term/phrase associations. FIG. 6A illustrates an index 610 in which a document/domain identifier 620 may be used as a primary key to look up one or more associated terms/phrases 630. The entries of the index 610 may be ordered using the document/domain identifiers 620.

[0061] FIG. 6B illustrates an alternative index 650 in which a document/domain identifier 660 may be used as a primary key to look up one or more associated {term, term count} and/or {phrase, phrase count} pairs 670. The pairs 670 may be ordered based on the counts. In this way, terms/phrases that have been used more frequently in queries that generated a search result, the selection of which led to the document, can be ordered ahead of those used less frequently. The entries of the index 650 may be ordered using the document/domain identifiers 660.

§ 4.2.2 Using Query Term to Selected Document Associations

[0062] FIG. 7 is a bubble chart illustrating operations that may use selected document to query term associations to suggest targeting keywords for an ad, or a set of ads. Query term/phrase to ad association operations 710 may use query term/phrase to selected document/domain association information 450 and ad information 720 to generate targeting keyword suggestions 730 for one or more ads. In an exemplary embodiment of the present invention, the ad information 720 may include one or more of an ad identifier, an ad creative, a landing page (e.g., a document identifier, such as a Web page URL), a Website home page (domain), an offer price (e.g., a bid or maximum cost offer (e.g., per impression, selection, conversion, etc.) for one or more targeting criteria, such as one or more targeting keywords, etc.), and search constraints (e.g., targeting keywords, geotargeting information, time/date targeting information, etc.) 725. In an exemplary embodiment of the present invention, the targeting keyword suggestions 730 may include one or more elements, each including an ad identifier and one or more suggested targeting keywords 735.

[0063] FIG. 8 is a flow diagram of an exemplary method 800 that may be used to associate query terms/phrases with advertisements, and use such query terms/phrases as one or more types of ad information. Document/domain identifier information is accepted. (Block 810) For example, such information may be accepted from ad information. More specifically, an advertiser may have specified a landing page URL and/or a Website URL for an ad. The accepted document/domain identifier may then be used to lookup associated terms/phrase. (Block 820) Recall, for example, exemplary data structures 610 and 650 of FIGS. 6A and 6B, respectively, in which a document/domain identifier 620/660 can be used as a lookup key to find associated terms/phrases 630, 670.

[0064] Having found associated terms/phrases, one or more of acts 830, 840 and 850 may be performed, depending on the desired application. For example, at least some of the terms/phrases may be imported as targeting keyword(s) for ad(s) associated with the document/domain. (Block 830) At least some of the terms/phrases may be provided to a user (e.g., an advertiser) as candidate targeting keywords for an ad(s) associated with the identified document/domain. (Block 840) At least some of the terms/phrases may be used to generate an (e.g., a candidate) ad creative. (Block 850) For example, the terms/phrases may be used to populate certain parts of a generic creative template.

§ 4.2.3 Refinements and Alternatives

[0065] Recall from FIG. 6B that counts may be associated with terms/phrases, which are associated with a document identifier. These counts may be used in conjunction with absolute and/or relative thresholds or tests in certain applications. For example, if the terms/phrases are to be provided as targeting keywords, they may have to have a certain minimum count to be considered. As another example, only the N terms/phrases with the highest counts might be considered. As yet another example, if the terms/phrases are to be provided as elements of an ad creative, only the term or phrase with the highest count might be used.

[0066] Such thresholding permits a number of useful features. For example, it may be desirable to ensure that a term or phrase will likely generate a number of impressions deemed sufficient. Alternatively, or in addition, thresholding can be applied to queries. For example, if a query occurred only one time (or only a very few times), or was issued by only one user, it may be discarded to preserve user privacy.

[0067] FIG. 9 is a flow diagram of an exemplary method 900 that may be used, at one or more places, to reduce the amount of data being stored. At decision block 910, it is determined whether or not a document/domain is associated with an ad or ads. If not, the document/domain information to query information association may be discarded. (Block 920) For example, referring back to FIG. 4, associated document information and query information 430 that has no associated ad may be discarded, or not generated in the first place. Referring back to FIG. 9, at decision block 930, it is determined whether or not an ad already uses a term or phrase as a targeting keyword. If so, the query term or phrase may be discarded. (Block 940) This can be used to avoid the need to save a suggested targeting keyword for an ad that is already being used by the ad. At decision block 950, it is determined whether or not an advertiser has previously rejected a term or phrase as a targeting keyword. If so, the query term or phrase may be discarded. (Block 940) This can be used to avoid the need to save a suggested targeting keyword that has already been considered and declined by an advertiser. Various other data filtering acts are possible and will often depend on the application(s) for which the data will ultimately be used.

[0068] As discussed above, terms/phrases can be associated with a selected document/domain. One application is to track keywords by document and to automatically generate a creative relevant to the document to association with such keywords. For example, many Websites may have Web pages with unique, non-overlapping content, and linking an ad to a landing page with content a user is interested in,